REMARKS

Claims 1-42 are pending in the present application. No claims were canceled, amended, or added. Reconsideration of the claims is respectfully requested.

I. 35 U.S.C. § 112, Second Paragraph: Claims 1, 15, 29, and 33

The Office Action rejects claims 1, 15, 29, and 33 under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which Applicants regard as the invention. This rejection is respectfully traversed.

With regards to claims 1, 15, 29, and 33, the Office Action states:

> 5. The following claim language is indefinite:
>
> a. With respect to the newly added limitation in the independent claims of "specifying a particular level of lag", it has not been demonstrated or made clear exactly where the "specifying" comes from, i.e. whether the "specifying" is done by a user or automatically, either statically or dynamically in response to changing conditions. The limitation will be given its broadest reasonable interpretation, i.e. any specification of a lag will suffice.
>
> b. It is unclear what the "synchronicity setting" is specified as. The setting could be a lag based on units of time or by workload. Applicant has failed to demonstrate a sufficient level of detail regarding the condition that a "synchronicity setting" is intended to regulate, i.e. time or workload. Claims 6 and 7 indicate that the synchronicity setting could be either a number of operations or a unit of time. Therefore, the claim will be treated as though specifying a level of lag in accordance with either a lapse of time or workload meets the claim limitations.

Final Office Action dated October 31, 2005, pages 2-3

The Office Action alleges that the term "specifying" is indefinite because "it has not been demonstrated or made clear exactly where the 'specifying' comes from, i.e. whether the 'specifying' is done by a user or automatically, either statically or dynamically in response to changing conditions." Applicants respectfully disagree. Applicants submit that the term "specifying" is not indefinite and that the comments presented in the Office Action point to a question of scope or breadth of the term "specifying," and not to the indefiniteness of the term.

In addition, the Office Action alleges that the term "synchronicity setting" is indefinite because "The setting could be a lag based on units of time or by workload. Applicant has failed to demonstrate a sufficient level of detail regarding the condition that a 'synchronicity setting' is intended to regulate, i.e. time or workload." Applicants respectfully disagree. The Office Action admits that claims 6 and 7 claim two separate exemplary embodiments wherein the "synchronicity setting" is a number of operations in one exemplary embodiment and a unit of time in the other exemplary embodiment. Applicants respectfully submit that the term "synchronicity setting" is not indefinite and that the comments presented in the Office Action point to a question of scope or breadth of the term "synchronicity setting," and not to the indefiniteness of the term.

A broad claim, no matter how broad, is not indefinite so long as the boundaries of the claim are capable of being understood. In other words, if the metes and bounds of a claimed invention are clearly ascertainable, then the claim, no matter how broad cannot be properly rejected as vague and indefinite" under the language of 35 U.S.C. § 112, second paragraph. *In re Gardner*, 427 F.2d 786, 166 U.S.P.Q. 138 (C.C.P.A. 1970) and *In re Goffe*, 576 F.2d 1393, 188 U.S.P.Q. 131 (C.C.P.A. 1975

Therefore the rejection of claims 1, 15, 29, and 33 under 35 U.S.C. § 112, second paragraph has been overcome.

II. 35 U.S.C. § 103, Obviousness: Claims 1-42

The Office Action rejects claims 1-42 under 35 U.S.C. § 103(a) as being unpatentable over Sicola et al., System for Data Replication Using Redundant Pairs of Storage Controllers, Fibre Channel Fabrics and Links Therebetween, U.S. Patent No. 6,601,187, July 29, 2003 (hereinafter "Sicola"). This rejection is respectfully traversed.

With regards to claim 1, the Office Action states:

> 7. As per claim 1, Sicola teaches the invention as claimed, including a method for synchronizing transactions, comprising:
> specifying a particular level of lag, said particular level of lag being a specified synchronicity setting (col. 11 lines 17-24; col. 12 lines 4-15);
> executing a series of commands at a first computing entity (col. 12 lines 17-25);

controlling a level of lag between computing entities by relaying the series of commands to a second computing entity (col. 12 lines 47-56) until said synchronicity setting is reached (col. 11 lines 17-24); and

wherein the second computing entity lags behind the first computing entity by an amount of lag that is no greater than said specified synchronicity setting (col. 11 lines 17-19).

8. Though Sicola does not specifically require postponing relaying additional commands after said synchronicity setting is reached, such is an obvious -modification of the asynchronous data replication method disclosed therein, particularly when viewed in relation to Sicola' s method of synchronous data replication. Sicola teaches that in a synchronous mode of operation, each input/output command that is entered at the host computer is replicated at a remote computer, and a subsequent command does not begin until the first command has been verified as complete (col. 11 line 43 - col. 12 line 3). However, certain constraints make this type of operation undesirable, as a system experiences downtime while waiting for the remote acknowledgement (col. 11 lines 35-42).

In contrast, asynchronous operation allows a series of commands to be entered at a first computer before requiring acknowledgment of completion of the replication at the remote site. This results in a lag between the operations, but is assured of eventual synchronization over time (col. 11 lines 17-24; col. 12 lines 6-15). This leaves open the question of what to do in the case of a delay at the remote site, which could be caused by numerous factors, including a simple bottleneck. Sicola is deliberately silent on this issue, as it is not within the scope of the disclosure (col. 12 lines 64-67, "If the remote copy was unsuccessful for other reasons, then... other error recovery procedures... are invoked.") Thus, there is a glaring need to fill in the blank of how to handle such a condition.

Final Office Action dated October 31, 2005, pages 3-4

The Office Action bears the burden of establishing a *prima facie* case of obviousness based on the prior art when rejecting claims under 35 U.S.C. § 103. *In re Fritch*, 972 F.2d 1260, 23 U.S.P.Q.2d 1780 (Fed. Cir. 1992).

Independent claim 1, which is representative of independent claims 15, 29, and 33 with regard to similarly recited subject matter, recites:

1. A method for synchronizing transactions comprising:

specifying a particular level of lag, said particular level of lag being a specified synchronicity setting;

executing a series of commands at a first computing entity;

> controlling a level of lag between computing entities by relaying the series of commands to a second computing entity until said synchronicity setting is reached; and
>
> postponing relaying additional commands after said synchronicity setting is reached, wherein the second computing entity lags behind the first computing entity by an amount of lag that is no greater than said specified synchronicity setting.

In comparing Sicola to the claimed invention, the claim limitations of the presently claimed invention may not be ignored in an obviousness determination. Claim 1 of the present invention recites the feature of "specifying a particular level of lag, said particular level of lag being a specified synchronicity setting." Such a feature is not taught or suggested by Sicola. Therefore, claim 1 is not obvious in view of Sicola because the features believed to be disclosed by this cited reference are not present.

The Office Action points to column 11, lines 17 through 24 and column 12, lines 4 through 15, reproduced below for the Examiner's convenience, as teaching this feature:

> When system **100** is in asynchronous mode, the remote site may lag behind by a bounded number of write I/O operations. All commands that are returned to the host as completed, are completed on the initiator, and may or may not be completed on the target. From a recovery viewpoint the only difference between the operation modes is the level of currency of target members.
>
> **FIG. 8A** is a flowchart showing asynchronous operation the present system **100**. Asynchronous operation provides command completion to the host after the data is safe on the initiating controller, and prior to completion of the target command. During system operation, incoming host write requests may exceed the rate at which remote copies to the target can be performed. Copies therefore can be temporarily out of synchronization, but over time that data will converge to the same at all sites. Asynchronous operation is useful when transferring large amounts of data, such as during data center migrations or consolidations.

The above cited passages of Sicola do not teach specifying a particular level of lag. The first cited passage of Sicola, column 11, lines 17 through 24, teaches a standard definition of the term "asynchronous mode." Sicola acknowledges that lag may exist in an asynchronous copy situation; however, Sicola does not specify that lag. The Office Action states:

> Sicola's synchronous mode of operation provides an obvious remedy that is the same as provided by the claimed invention, namely to wait for acknowledgment of completion from the remote site. It could be said that Sicola inherently discloses this feature, or at the least that it obviously follows from the open-ended teachings of the asynchronous mode of operation. Specifically, Sicola states that "the remote site may lag behind by a bounded number of write 110 operations." (col. 11 lines 17-24). Thus, when the number of operations the remote site lags behind the host reaches the bound, it would have been obvious to one having ordinary skill in the art to cease initiation of new operations until at least one completion notification has been received. The discussion presented herein relates to all other independent claims presented in this application, and is hereby incorporated by reference into the rejections of those claims.

Final Office Action dated October 31, 2005, page 5

It appears that the Office Action is equating the phrase "bounded number of write I/O operations" to a specified synchronicity setting, as recited in the present invention. Applicants submit that the phrase "bounded number of write I/O operations" is not the same as a specified synchronicity setting and that the only way to reach this conclusion is through an impermissible use of hindsight with the benefit of applicants' invention as a model. Sicola does not specifically define what bounded means. However, Sicola uses the term "bound' in two other places, column 8, line 65 and column 10, line 66. In both passages the term implies a meaning of connected or grouped. Therefore, one valid interpretation is that the term "bounded number of write I/O operations" means a number of connected or grouped write I/O operations. Thus, Sicola is merely stating that the lag could be more than one write I/O operation in length. Furthermore, the write I/O operations could be bound by any number of causes, such as the write/execution speed of the target or the storage capacity of the target, etc. The only way to interpret "bounded number of write I/O operations" as meaning a specified synchronicity setting is through an impermissible use of hindsight with the benefit of applicants' invention as a model.

Acknowledging that lag exists in an asynchronous mode is not the same as specifying a particular level of lag. Sicola does not teach specifying a particular level of lag. Furthermore, in column 12, lines 4 through 15, reproduced above, Sicola teaches away from the feature of "specifying a particular level of lag, said particular level of lag being a specified synchronicity setting." In the above cited passage, Sicola teaches that over time, an asynchronous system will

become synchronous, if allowed to operate long enough. Specifically, Sicola states "Copies therefore can be temporarily out of synchronization, but over time that data will converge to the same at all sites." (emphasis added) As Sicola teaches that lag dissipates over time, Sicola cannot teach specifying a level synchronicity for the system to operate at or controlling lag between computers by maintaining a synchronicity setting. Therefore, Sicola does not teach or suggest the feature of "specifying a particular level of lag, said particular level of lag being a specified synchronicity setting," as recited in claim 1 of the present invention. Thus, The Office Action fails to state a *prima facie* case of obviousness.

Additionally, claim1 recites the feature of "controlling a level of lag between computing entities by relaying the series of commands to a second computing entity until said synchronicity setting is reached." Sicola does not teach or suggest this feature. Sicola does not teach controlling a level of lag. As discussed above, Sicola merely acknowledges that lag may exist. Sicola does not provide for controlling a level of lag. Because Sicola does not teach controlling a level of lag, Sicola does not anticipate Applicants' claims.

Further, Sicola does not teach controlling a level of lag by relaying commands until the synchronicity setting is reached. Nothing in Sicola teaches controlling lag by relaying commands until a synchronicity setting, which is the specified particular level of lag, is reached. The Office Action alleges that column 12, lines 47 through 56, of Sicola, teaches controlling a level of lag between computing entities by relaying the series of commands to a second computing entity. The Office Action alleges that column 11, lines 17 through 24, reproduced above, teaches "until said synchronicity setting is reached." Sicola, column 12, lines 47 through 56 states:

> At step **830**, PPRC manager **515** (via host port initiator module **510**) sends the write data to the remote target. Order preserving context is also passed to host port initiator module **510**. At step **835**, remote target controller B1 (**211**) writes data to its write-back cache (or associated media if a write-through operation). A check is then made by controller A1 at step **840** to determine whether the remote copy successfully completed. If so, then, at step **845**, target controller B1 sends the completion status back to initiator controller A1.

The above cited passage of Sicola does not teach controlling a level of lag between computing entities by relaying the series of commands to a second computing entity. Instead, the above cited passage of Sicola merely teaches the normal asynchronous operation of system, wherein

controller A1 checks to see if a write operation has been completed by the remote target. The above cited passage does not teach or suggest anything about lag, or controlling a level of lag. The passage only teaches two systems communicating with each other. As discussed above, column 11, lines 17 through 24 does not teach a synchronicity setting or reaching a synchronicity setting. Column 11, lines 17 through 24 merely teaches that lag exists in an asynchronous system. Because Sicola does not teach controlling a level of lag by relaying commands until a synchronicity setting is reached, Sicola does not anticipate Applicants' claims.

Furthermore, claim 1 recites the feature of "postponing relaying additional commands after the synchronicity setting is reached, wherein the second computing entity lags behind the first computing entity by an amount of lag that is no greater than the specified synchronicity setting." The Office Action admits, and Applicants agree that Sicola does not teach postponing relaying additional commands after the synchronicity setting is reached. However, the Office Action states that the feature is an obvious modification to Sicola. Applicants respectfully disagree. The Office Action states:

> This leaves open the question of what to do in the case of a delay at the remote site, which could be caused by numerous factors, including a simple bottleneck. Sicola is deliberately silent on this issue, as it is not within the scope of the disclosure (col. 12 lines 64-67, "If the remote copy was unsuccessful for other reasons, then... other error recovery procedures... are invoked.") Thus, there is a glaring need to fill in the blank of how to handle such a condition.

Final Office Action dated October 31, 2005, page 4

Thus, the Office Action is stating that since there is an obvious, glaring problem in the Prior Art that needed to be solved, and the Applicants solved the obvious problem, the Applicants' solution must be obvious. This is not proper logic. Some time after the chariot was invented, it must have become obvious that a way was needed to stop the vehicle besides hoping that the horse would stop. Thus, someone invented the first brake system. However, the invention of the brake was not rendered "obvious" as that term is defined for 35 U.S.C. § 103(a), simply because the need for it was obvious or notorious. Furthermore, Sicola points to the fact that the solution arrived at by the present invention is not obvious. In column 11, lines 13 through 14,

Sicola states that "steady state operation is possible in two modes, synchronous and asynchronous." Sicola further states that asynchronous operation is assured of eventual synchronization over time (col. 11, lines 17-24; col. 12, lines 6-15). Neither of these statements points to a solution of imposing constraints upon an asynchronous operation so that it maintains a steady, synchronous state of asynchronization.

Additionally, modern gaming systems today operate in either a synchronous or asynchronous state, depending upon which is best suited for the application; there is no combination of the two. First person shooter type games usually use asynchronous mode because it is not necessary to keep precise track of everything, whereas sports games, like football, often use synchronous mode as it is necessary to keep track of where everyone or everything on the board or field is.

As it appears that the Office Action is asserting that the modification would be obvious to one of ordinary skill in the art under M.P.E.P. § 2144.03, Applicants respectfully challenge this assertion and request that the Examiner provide the evidence necessary to show that such a modification would be obvious. Otherwise, the Office Action has merely made a modification to Sicola without any basis in the prior art and has failed to establish a *prima facie* case of obviousness. If the rejection is based on personal knowledge, Applicants respectfully request that the Examiner provide an appropriate affidavit such that the Applicants may review the affidavit and provide responding affidavits to contradict or explain if appropriate.

Therefore, for all the reasons set forth above, Applicants respectfully submit that Sicola does not teach or suggest the features of claims 1, 15, 29, and 33. At least by virtue of their dependency on claims 1, 15, 29, and 33 respectively, Sicola does not teach or suggest the features of dependent claims 2 – 14, 16 – 28, 30 - 32, and 34 - 42. Accordingly, Applicants respectfully request the withdrawal of the rejection of claims 1-42 under 35 U.S.C. § 103(a).

In addition, Sicola does not teach or suggest the specific features of dependent claims 2 – 14, 16 – 28, 30 - 32, and 34 - 42. For example, with regard to claims 8, 9, and 32, the Examiner stated that Sicola teaches the amount of lag and the specified synchronicity setting being measured as amounts of data or number of devices. The Examiner takes Official Notice that there are well known ways of specifying lag such as through the use of buffers.

Sicola, however, does not teach specifying lag. Sicola states that lag exists but does not teach a way of specifying the lag. Sicola does not teach controlling a level of lag. Sicola does

not teach controlling lag by relaying commands until the synchronicity setting is reached. Sicola does not teach postponing relaying additional commands after the synchronicity setting is reached. Therefore, the combination of Sicola with the use of buffers does not render Applicants' claims unpatentable.

Sicola does not describe, teach, or suggest specifying a particular level of lag, the particular level of lag being a specified synchronicity setting, a level of lag between computing entities being controlled by relaying commands until the synchronicity setting is reached, or the relaying of additional commands being postponed after the synchronicity setting is reached. Therefore, Sicola does not anticipate Applicants' claims or render them unpatentable.

III. Conclusion

It is respectfully urged that the subject application is patentable over the cited references and is now in condition for allowance.

The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

DATE: Dec. 19, 2005

Respectfully submitted,

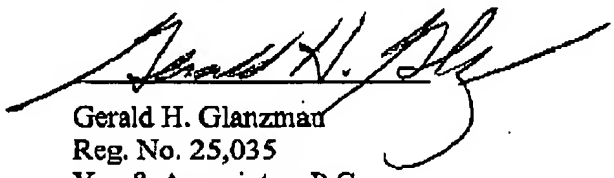

Gerald H. Glanzman
Reg. No. 25,035
Yee & Associates, P.C.
P.O. Box 802333
Dallas, TX 75380
(972) 385-8777
Attorney for Applicants

GHG/BJ